

# Enhancing Programming Interface to Effectively Meet Multiple Information Needs of Developers

Haipeng Cai<sup>1</sup>

Department of Computer Science and Engineering, University of Notre Dame  
Notre Dame, IN, 46556, USA  
hcai@nd.edu

**Abstract.** In the past decades, integrated development environments (IDEs) have been largely advanced to facilitate common software engineering tasks. Yet, with growing information needs driven by increasing complexity in developing modern high-quality software, developers often need to switch among multiple user interfaces, even across different applications, in their development process, which breaks their mental workflow thus tends to adversely affect their working efficiency and productivity.

This position paper discusses challenges faced by current IDE designs mainly from working context transitions of developers during the process of seeking multiple information needs for their development tasks. It remarks the primary blockades behind and initially explores some high-level design considerations for overcoming such challenges in the next-generation IDEs. Specifically, a few design enhancements on top of modern IDEs are envisioned, attempting to reduce the overheads of frequent context switching commonly seen in the multitasking of developers.

**Keywords:** Information need, integrated development environment, context switching, automatic recommendation, programming interface, software visualization

## 1. Introduction

One merit with visual programming [3, 6, 13] is that its integrated interface empowers smooth transitions among the workflow steps of developers—the interface provides all programming elements (of visual forms) so that the developers involved in the interface can easily maintain their mental workflow models by focusing on mostly just one type of interface (i.e., visual).

With most existing IDEs (e.g., ECLIPSE [1]), however, developers often face challenges from frequent transitions between coding (text interface) and visual aids (graphical interface), or even between disparate applications (and their different interfaces) [12]. Since traditional (textual) programming involves typically a demanding logic reasoning process, such transitions and context switches can cause great inefficiency [5] in the development activity, even greater risks to the quality of resulting software.

The reason underneath is that context switching tends to interrupt the workflow [5] of developers. More important, this problem can be even exacerbated by the growing information needs for developing modern software of increasing scale and complexity. Unfor-

tunately, on the other hand, modern IDEs tend to grow in the complexity of their interface in a way that, seemingly facilitating developers to meet their needs for multiple sources of information, actually compounds the problems with switching among increasingly more contexts.

As it stands, research on easing programming tasks through interactive graphical environments exists [2] with most focusing on providing visual aids within IDEs. For instance, Dragon [4] shows visual windows for program dependence, debugging information, data structure state, memory layout and similar other visual gadgets during program analysis tasks of developments, incorporated into the entire software development work flow. However, this framework is limited to passively responding to user requests—it fails to automatically *push* information to assist program-analysis tasks as if it were an integral part of the full task pipeline. Thus, a more useful interactive programming environment needs to deliver informative visual aids not only on demand but in a proactive manner, so as to minimize context switches during the entire development workflow.

Another challenge to today’s IDE interfaces lies at their falling short of meeting the growing amount and *variety* of information needs by developers. To finish a coding task, for example, developers usually have to consult many information sources that are diverse and distributed across disparate interfaces even applications. Although modern IDEs mostly have strong supports to integrate diverse functionalities by means of plug-ins or extensions, information from those extraneous modules often has to be passively retrieved under the requests of developers—the multiple sources of information are available, yet not well integrated *in synergy* [15] with other elements of the IDEs such that developers can concentrate on their holistic workflow.

To address these specific issues, this paper preliminarily explores several novel IDE features that could effectively assist developers with handling multiple tasks while minimizing the costs of context switching during their development workflow. Specifically, three features are proposed focusing on interface design: extending traditional coding view to include co-worker views, offering automatic recommendation based information for API usage and code examples, and providing in-situ mechanism for mostly common used code-editing related operations driven by current the task context. And two major visualization features are discussed, including multiple code visualization views and interactive linked visualizations. By illustrating the needs and benefits of these instrumental features, the paper demonstrates how the next-generation IDEs could be designed to offer better aids to developers in ways that improve development efficiency and productivity.

In summary, this *position* paper highlights the context-switching issue in the design of today’s IDEs that hinders the effectiveness of using them, and illustrate such issue using example usage scenarios; it discusses three interface design features that potentially reduce developers’ overall cost of switching among multiple contexts in search of various sources of information; it also envisions two interactive visualization features that enable holistic integration of multiple information in synergy so as to reduce developers’ need of switching contexts when searching for various information.

The rest of this paper is organized as follows. First, Section 2 gives a development scenario regarding information foraging that motivates our programming interface design. Then, Section 3 and Section 4 summarize the concrete features in the new programming environ-

ments, on interface and visualization design, respectively. Finally, Section 5 outlines the next step, planning on the implementation and evaluation of the proposed design.

## 2. Motivating Example

During software development, programmers gain most of the information they need from the source code they are working on. Yet, they also need information beyond that, such as those produced by program analysis tools, to obtain better understanding of the software. Examples of such additional information include call graphs, program dependencies, and type hierarchies. While most present IDEs do provide functionalities to help developers obtain these information, they force developers to actively make requests for them. However, responding to user requests may not be sufficient in many situations. Rather, a more effective IDE should provide developers with a *guiding* interface instead of question responder, as developers may not have *prerequisite* information for them to initiate those requests or to do so in the most efficient way overall. In consequence, excessive context switches ensue when developers have to resort to other contexts or even applications for obtaining those missing information.

In a typical usage scenario, a developer wants to know the overall design of the component-level architecture of a software for which he just finished the coding for one of its many packages. With a program analysis tool integrated in the IDE he is using, the developer *has to* choose a button or menu item relevant to the functionality on the call graph of the entire program. Further, the developer proceeds by looking for all possible interfaces compatible for a function call of interest. Thus the developer has to traverse the call graph and hover mouse cursor over all relevant modules one by one.

However, without rich experiences with the very details of this software, it is infeasible for the developer to know how to make the preceding requests. The key issue, which is really the main obstacle here, is the requirement for the user to recognize which requests to make without auxiliary information provided by the program analysis tool. As such, the value of this visual-aid tool itself apparently diminishes. Arguably there exists a crucial need of developers with an IDE that incorporates interactive program analysis tools is a workflow-driven pipeline where the transitions from graphical to textual settings, and of course the other way around, are as seamless as possible. We are motivated by such an observation and the consequent requirement in the design of interactive programming interfaces.

## 3. Interface Design

Developers spend most of their time on their code for adding new features, making changes, debugging, and code comprehension [11]. When doing these tasks, developers often need also external assistances integrated in their workflow (e.g., automatic code completion [14]) which facilitate their development efficiency. To meet such needs, a tentative framework could incorporate three interface design features to further reduce workflow transitions of developers when they are working around their code.

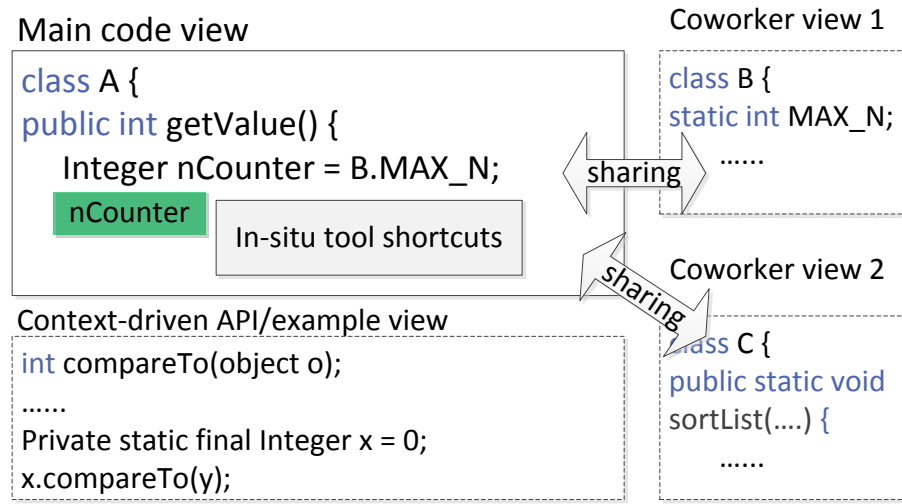


Fig. 1 A new interface design that helps reduce context switches of developers between coding and getting aids, and supports close collaborations among coworkers in software development teams.

Figure 1 gives an overview of these design features. Aside the traditional coding view, there are a few other co-worker views that assist with communication and collaboration tasks typical in team development scenarios; at the bottom, the context-driven API/example view attempts to provide code examples that are recommended based on current coding context to assists programmers with using APIs of which usages are not familiar to them; finally, the in-situ interface shown in the main code view illustrates the design of porting convenience invocation shortcuts, which are mostly spread over varying places in existing IDEs, to the current focus of editing.

### 3.1. Context-driven API/example View

While coding, programmers often have questions about the usage of some third-party functionalities or features [9]. And while implementing a feature, they face hard questions concerning which functions or objects they should pick [10]. To some extent, these questions can be reduced to the needs for getting function usage information and, even further, illustrations of that usage with example code. Active code completion [14] via API menus already helps developers better than using separate API browsing views, yet it may not be sufficient as developers have to navigate through possibly long API lists (and hovering on each one to get the function prototype or API documentation on a floating window as seen in ECLIPSE [1]), which could potentially break their mental model focusing on the programming logic.

It is plausible to, optionally, put such assistance back into a separate view but closely connected to the main code editing view (as shown at the bottom of Figure 1), where usage information of relevant APIs is display on demand based on the current context of

object accesses or function calls. Importantly, all relevant APIs are ranked according to their frequency of being used recently at the default mode. Similar solutions have actually been explored previously in a more general sense from a perspective of the information foraging theory, with respect to software engineering tasks such as programming and debugging [12].

A more important reason for providing the option of moving API usage information to a separate view is the need of combining code examples with the usage. While showing function prototype and/or API documentation is helpful to developers to fill in arguments, it is more beneficial to show them usage examples thereof with the usage synopsis. In practice, programmers search code examples with respect to unfamiliar APIs very often, by using Internet searches, for instance, even preferably over reading API documents.

In this regard, at least three sources of search for such code examples can be taken into account. The first one is the examples coming with APIs in their documentation. Another option is searching in the current code base for relevant examples using context similarity measurement (e.g., calling context and/or type of the object from which the API would be invoked). The code shown in the API/example view of Figure 1 illustrates the result obtained from this source: When the cursor lies immediately after the `Integer` object `nCounter`, the view shows candidate API lists applicable to objects of the `Integer` type, with ones most frequently used recently listed at the top (`compareTo` here), and followed by the code example found in the current code base. Such examples give an instant and clear demonstration on how to use the relevant APIs. Finally, an automatic web search, using open search engine programming interface (e.g., Google API), can be initiated with queries concerning the function usage (e.g., “`strtol` C++ example”). Then relevant content can be extracted and put back to the API/example view for programmers’ reference.

### 3.2. Coworker Views

Another key interface design feature of our framework concerns about the information needs of developers collaborating in a development team. Previous studies show that in collaborative development one of the primary information sources for developers is their co-workers [8]. In fact, it is very common that when developers have questions regarding how a function or feature is implemented, they tend to first resort to their teammate instead of software documentations [11]. To facilitate developers to take advantages of having co-workers to consult as their information needs arise, it is potentially rewarding to incorporate a set of coworker views aside the main code editing view (shown on the right-hand side of Figure 1).

The rationale of introducing these additional views is two-fold. First, developers working in the same team can easily share their source code in real-time when necessary. One example case in which this sharing could be useful is when a senior developer coaches a team member in familiarizing him with the team project. Another example can be seen in agile development, where one developer could quickly prototype his function according to the ongoing implementation of a function being written or debugged by

another developer. As shown in Figure 1, the current developer is coding the method `getValue()` for class A, with a reference to the static variable `MAX_N` of class B that is being coded by a coworker. Having the choice of checking the implementation of a component, developed concurrently by a teammate, on which current coding task is dependent will save a developer's time seeking for the information about that component in other more expensive ways.

Second, such views can enable close collaboration among physically distributed teammates. For instance, a developer who needs one of his teammates to demonstrate how to write or debug a piece of code would readily get the help from such views without moving to a different seat or office, or resorting to another instant messaging tools. Screen space allowing, such benefits can be even augmented with multiple coworker views are opened at the same time—allowing close collaborations among multiple developers.

At the first glance, the above interface designs seemingly conflict our goal of reducing context switches by developers, because those extra views potentially end up with more context switches. However, the overall cost of context switching will be mostly reduced indeed as the total time developers would spent on getting the information from these views can be much greater without these integrated views and information. Consider finding the code example for an API again. Without the automatic code reference shown within the IDE, a developer would have to search online or consult to other sources that are available usually in different interfaces from the whole IDE (e.g., a different application like web browser).

### 3.3. In-situ Interface Elements

Almost all IDEs today contain a main menu at the top of the entire interface, followed by one or several rows of tool shortcuts shown as buttons or icons. Although usually those menus or shortcuts can be situated differently, few of them is tightly incorporated into the working area of developers where their functionalities will be applied to. For example, there is always a considerable “visual distance” from the code being focused on by developers and the shortcuts to functions developers need to utilize on that code. While the context switches in such situations are not as large as those seen in cases where developers seek coworker resources without coworker views, such distance could be much reduced. Accordingly, two possible interface improvements to reduce the unnecessary distance can be investigated.

First, in-situ tool shortcuts can be added to the main code editing view. The presence of such gadgets is contingent on user actions of marking focus on (e.g., selecting) code elements to which the shortcuts are applicable; and the composition of the gadgets is determined by the characteristics of the code elements being focused on by developers. As an example, Figure 1 shows, in the main coding view, an “in-situ tool shortcut bar” aside the object `nCounter` when it is selected through double-click and the mouse cursor hovers nearby—the gadget disappears once the selection is revoked or the cursor moves away the focused object. This is akin to the in-situ formatting toolbar in Microsoft Office, triggered by double-clicking on a word.

The more important part of this design is the demand-driven composition of the gadget. To effectively reduce perceptual transitions within the IDE, the in-situ tool gadget should contain most, if not all, shortcuts to functionalities that developers would possibly use for the focused object. This decision can be made in reference to developers' common information needs with respect to that object, based on such criteria as the object type. For instance, for a function identifier in its invocation statement, example shortcuts would be "caller list", "rename", "declaration" and so on.

Second, the presence and layout of visual components should be demand-driven. As developers usually work on multiple tasks during their development workflow [11], they tend to switch among multiple sources of information. Yet, they can mostly focus on one task or information at a time only. The IDE thus needs to optimize the size and composition of the particular visual space where a developer has to concentrate on for a particular task, while diminishing the presence or even phasing out visual components irrelevant to the current task.

For example, when a developer is right in the process of typing code, visual components, such as the top menu and main tool bar, side panels, and bottom debugging views, becomes irrelevant and thus should automatically disappear so that the main coding view gets its maximal visual space. Some IDEs, such as Microsoft Visual Studio, has already incorporated similar features (e.g., dockable gadgets), yet relies on user settings to apply those features. Also, the overall design there does not support automatic adaptation of interface layout and composition to user action and workflow contexts. The dockable gadgets, for instance, can be set to hide when mouse cursor moves away from them, but the layout and elements of the gadgets do not automatically accommodate developers' information needs varying in the development workflow. In addition, all visual components that are not applicable to the current developer action can phase out from the interface and come back when they become applicable again. In contrast, most IDEs choose to disable those components while leaving them in the visual space.

## **4. Visualization Design**

During software development, programmers gain most of the information they need from the source code they are working on [10, 11]. Yet, they also need information beyond that [9], such as those produced by program analysis tools, to obtain better understanding of the software [15]. Examples of such additional information include call graphs, dependence graphs, and type hierarchies. While most present IDEs do provide functionalities, via plug-ins for instance, to help developers obtain these information, they force developers to switch among different visualizations of those information, potentially leading to expensive workflow interruptions. An additional issue is that usually these visualizations are separate from each other, without explicit links among them, forcing developers to maintain an extra mental model linking those information in their mind when manipulating them. In this context, it is reasonable to leverage multiple linked visualizations of source code, along with the source code itself, to facilitate the code understanding and navigation for developers.

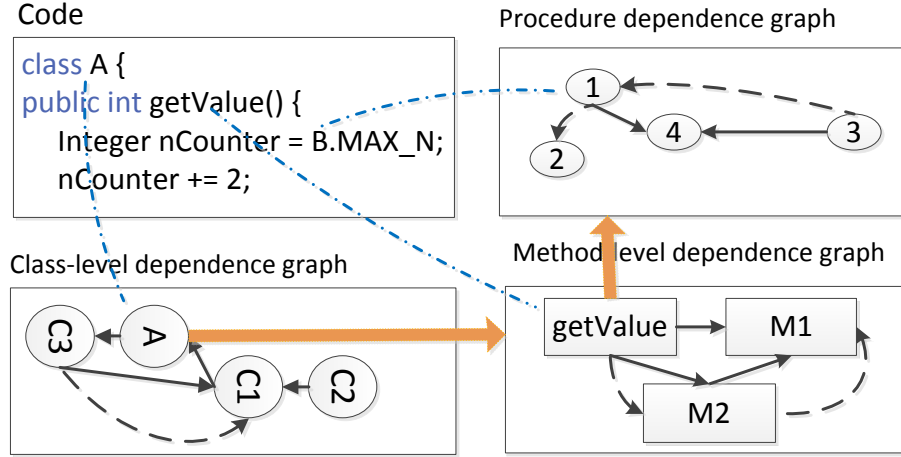


Fig. 2 Multiple linked visualizations of source code to integrate multiple code information in synergy.

#### 4.1. Multiple Visualizations of Source Code

Not only can information visualization greatly aid data understanding, but also can multiple forms of information of the same data set be even more helpful (e.g., [7]). During their mental workflow for code understanding, which is their primary task [11], developers can greatly benefit from multiple visualizations of the source code information besides the textual source code itself.

A relevant proposal would be to utilize multiple visualizations of program source code, which are interconnected underneath the source code, to enable more effective program understanding. Figure 2 illustrates our visualization design feature for the next-generation IDEs using dependence graphs as the example information representation of source code (there are other forms of such information of code, such as call graphs and type hierarchies as mentioned above). Surrounding the traditional code editing view (upper left) that provides the textual representation of source code only, three other satellite views show the dependence graph of the source code at three levels of detail, namely the (statement-level) procedure dependence graph (PDG) (bottom left), method-level dependence graph (MDG) (bottom right), and class-level dependence graph (CDG) (upper right). Different styles of arrows in the graphs illustrate different types of dependencies. The three dotted lines across the views illustrate the links between the source code and each of these graphs.

Optionally, these visualizations can be selectively or all added to the IDE. A more synergistic design is to *seamless synthesize* the four views *in one*. In the latter design, instead of showing more than one or all views at the same time, only one view is visible at a time. The motivation is that, again, while developers can be greatly benefited from multiple visualizations, they could utilize one of them at one time only. The idea is to switch among these visualizations by *zooming in/out* operations (shown by two wide arrows in



the figure). As with Google Maps, when developers zoom out from a statement, they will first switch to the PDG visualization with central points automatically set to that statement; then they can navigate on the PDG and zoom in from any node thereof back to the source code view. If developers continue to zoom out when they are in the PDG visualization, they switch to the MDG view where they can also navigate, at the method level, and zoom in back to the lower levels of details. Switching between MDG and CDG is similar. Alternatively, zooming in from a method declaration while in the source code view can directly lead the developer to the MDG visualization, and similarly, zooming in from a class declaration in the source code leads to the CDG directly.

#### **4.2. Interactions across Linked Visualizations**

The multiple visualizations of source code are also mutually linked, since they all representation the same data (source code). An additional merit of the multiple linked visualization is more effective interaction. One simple example is that selecting all lines of code of method can be more easily done by just selecting that method on the MDG visualization view; selecting or deleting a whole class will see greater benefit in similar ways. Moving code around through interactions on the dependence graph visualizations would be even more beneficial. For example, a developer can quickly start writing a method by cloning an existing one by copying the corresponding node on the MDG view. When multiple visualizations are shown simultaneously, more interactions can be enabled, such as moving or copying a method from one class to another. Of course, feasible interactions on the graphical visualizations are subject to feasible automatic source code level operations.

### **5. Conclusion and Future Work**

Today's developers usually deal with multiple tasks simultaneously during their software development process, seeking variously sources of information for interleaving tasks such as coding, documenting, testing, and debugging. To help them meet such needs, modern IDEs try to incorporate increasing number of interface elements to provide sources for meeting those multiple information needs, yet mostly are inclined to actually aggravate the problem of imposing on the developers the demands of frequently switching among many different contexts.

This paper thus explores in this regard and envisions several interface and interactive visualization design features for enhancing today's IDEs in a way that helps developers meet multiple information needs more efficiently. It illustrated the needs and benefits of incorporating those features in next-generation IDEs with motivating examples.

Beyond what has been tentatively proposed in this paper, there are potentially much more similar features than exemplified to be explored in the future, yet the discussions here could enlighten new lines of research improving programming interfaces and environments.

An immediate next step would be to develop relevant prototypes of the proposed features on the basis of a popular IDE such as ECLIPSE, and then to evaluate the feasibility and usefulness of such features through user studies with professional developers.

## References

1. Eclipse luna. <https://eclipse.org/>
2. Benton, W.C., Fischer, C.N.: Interactive, scalable, declarative program analysis: From prototype to implementation. In: Proceedings of the 9th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming. pp. 13–24. PPDP '07 (2007)
3. Cai, H., Chen, J., Auchus, A.P., Laidlaw, D.H.: Composing dti visualizations with end-user programming. CoRR abs/1310.2923 (2013)
4. Chapman, B., Hernandez, O., Huang, L., Weng, T.h., Liu, Z., Adhianto, L., Wen, Y.: Dragon: An open64-based interactive program analysis tool for large applications. In: International Conference on Parallel and Distributed Computing, Applications and Technologies. pp. 792–796 (2003)
5. Czerwinski, M., Horvitz, E., Wilhite, S.: A diary study of task switching and interruptions. In: Proceedings of the SIGCHI conference on Human factors in computing systems. pp. 175–182. ACM (2004)
6. Green, T.R.G., Petre, M.: Usability analysis of visual programming environments: a cognitive dimensions framework. *Journal of Visual Languages & Computing* 7(2), 131–174 (1996)
7. Hanciles, B., Shankararaman, V., Munoz, J.: Multiple representation for understanding data structures. *Computers & Education* 29(1), 1–11 (1997)
8. Ko, A.J., DeLine, R., Venolia, G.: Information needs in collocated software development teams. In: Proceedings of the 29th international conference on Software Engineering. pp. 344–353 (2007)
9. LaToza, T.D., Myers, B.A.: Developers ask reachability questions. In: Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering-Volume 1. pp. 185–194 (2010)
10. LaToza, T.D., Myers, B.A.: Hard-to-answer questions about code. In: Evaluation and Usability of Programming Languages and Tools. p. 8 (2010)
11. LaToza, T.D., Venolia, G., DeLine, R.: Maintaining mental models: a study of developer work habits. In: Proceedings of the 28th international conference on Software engineering. pp. 492–501 (2006)
12. Lawrance, J., Bogart, C., Burnett, M., Bellamy, R., Rector, K., Fleming, S.D.: How programmers debug, revisited: An information foraging theory perspective. *Software Engineering, IEEE Transactions on* 39(2), 197–215 (2013)
13. Metoyer, R., Lee, B., Riche, N., Czerwinski, M.: Understanding the verbal language and structure of end-user descriptions of data visualizations. *ACM SIGCHI Conference on Human Factors in Computing Systems* (2012)
14. Omar, C., Yoon, Y., LaToza, T.D., Myers, B.A.: Active code completion. In: Proceedings of the 2012 International Conference on Software Engineering. pp. 859–869. IEEE Press (2012)
15. Zeller, A.: The future of programming environments: Integration, synergy, and assistance. In: *FoSE*. pp. 316–325 (2007)